Santa Barbara Three-Dimensional (SB3D) Radiative Transfer Model

W. O'Hirok and C. Gautier
Institute for Computational Earth System Science
University of California, Santa Barbara, CA 93106

1. Introduction

The Santa Barbara three-dimensional (SB3D) radiative transfer model is a research tool used to understand the radiative environment in both the atmosphere and ocean for process and remote sensing studies. It is based on the Monte Carlo method. Put simply, a weighted photon enters the model domain, travels a specified distance that depends on the extinction coefficient of the atmospheric or oceanic constituent, is de-weighted by the single scattering albedo of the constituent, and is scattered along a trajectory defined by the phase function. This process is repeated until the photon exits the model domain or the weight of the photon is reduced to below a predefined threshold. After collecting statistics on a sufficient number of photons, fluxes can be estimated. Variance reduction methods have been added to enhance estimates of radiance. The spectral range of these computations is from the ultraviolet (0.25 μm) to the thermal (50.0 μm). In the thermal region backwards propagation is utilized. Constituents in the atmosphere include gases, aerosols, cloud droplets, and ice particles and in the ocean, hydrosols, chlorophyll and yellow matter. At the air-sea interface three-dimensional ocean waves are provided. Their angular distribution and foam content are computed as function of wind speed. The model can be run on any Windows-based PC.

2. Model Description

The model domain consists of a variable three-dimensional structure composed of cells that each represent a homogeneous atmospheric volume. Variability within the model is based on the spatial distribution of these cells. While the model has a maximum of 511 vertical layers, the only limitation on the number of cells in the horizontal direction is computer memory.

The basic computational steps in processing a photon are illustrated in Fig. 1. All photons are initially assigned a weight (see below) and placed at the top of the cell in the uppermost layer of the model domain. Microphysical properties for each constituent in that cell are retrieved from a pre-tabulated database (step 1). The distance a photon traverses between interactions within a cell is given by the pathlength, s, that is a function of a random number, R_s, and the extinction coefficient, κ_e, for a given wavelength, λ (step 2). The extinction coefficient is derived from the total optical depth of the cell, τ. If a photon escapes a cell, that part of the pathlength not used in the original cell is employed within the adjacent cell but modified to reflect the extinction coefficient of the new cell. The type of particle interaction a photon experiences within a cell is based on the ratio of a particle's optical depth, τ_i, to the total optical depth, τ, of the cell. A random number, R_i, is generated, and a particle is selected according to R_i's location with regard to the cumulative probabilities of each atmospheric constituent in that cell (step 3). The photon is then scattered the direction, θ, from its previous trajectory based on the probability distribution, $P_p(\theta)$, of the scatterer's normalized phase function $P(\theta)$ (step 4). During the interaction the energy absorbed is the weight of the photon, p_{wt}, multiplied by the single scattering albedo, ω_0, of the particle. The scattered photon's weight is reduced by the coalbedo, $1-\omega_0$ (step 5).

In the thermal region the processing of photons follows a similar path as demonstrated in Fig. 2 but uses a backwards propagation method. The initial location of the photon is at the position of the sensor and the photon's path is traced backwards towards the emitting source. The weight of the photon becomes a function of the emitter's temperature.

In its original version, SB3D optical properties for gaseous transmission were based on the 3-term k-distribution method from LOWTRAN7. More recently the gas properties have been supplemented with an 8 and 16 term k-distribution method derived from the HITRAN database (Yang et. al., in press). The spectral resolution is 5 nm from 0.25 to 1.00 μm and 10 nm beyond 1.00 μm. Gases include water vapor (H_2O), ozone (O_3), CO_2, N_2O, CO, CH_4, O_2, N_2, NH_3, HNO_3, NO, and SO_2. Computations of transmittance from a discrete-ordinates model show the new k-distribution

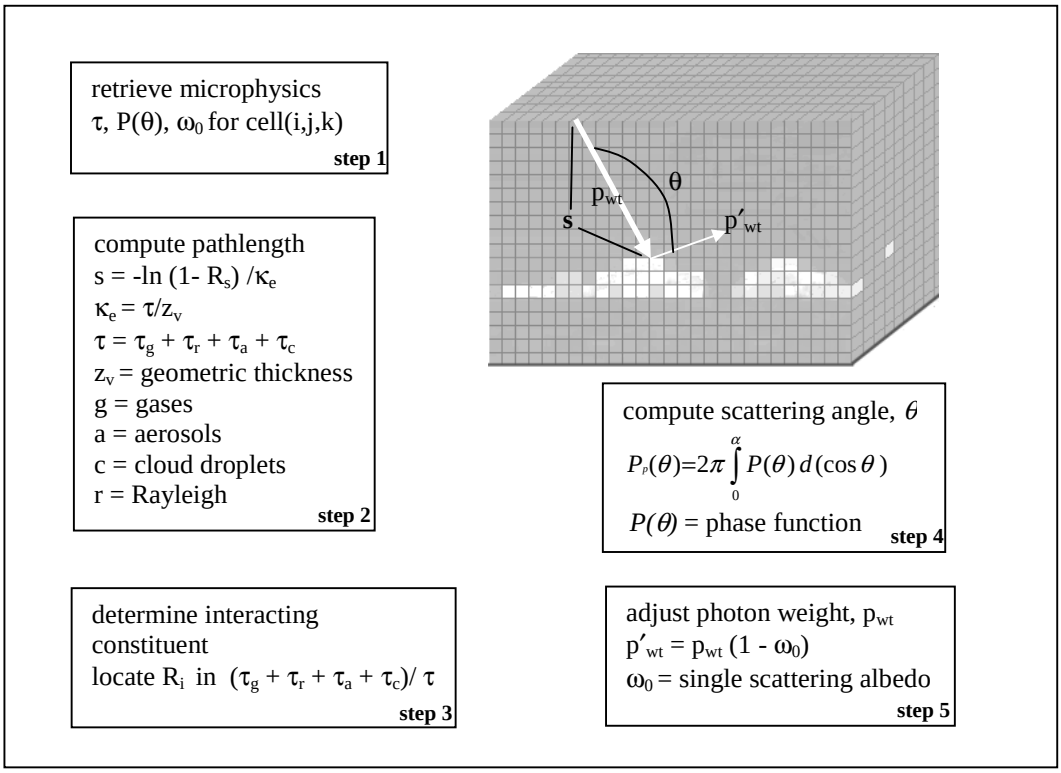

Fig. 1. SB3D computational steps.

method to be very close to those obtained from a line-by-line model. The number of k-terms used in the gaseous component translates directly into the number of times a photon must be processed in SB3D with the weight of the photon adjusted to each k-term, Δg weight. Although gaseous absorption by definition has a single scattering albedo equal to 0, ω_0 is actually set to a user-defined η to reduce the statistical variance in regions that may be surrounded by high gas concentrations. If η is set > 0, the gas extinction coefficient is rescaled by $(1 - \eta)^{-1}$. The direction of travel by a photon is not changed after a gas interaction.

Rayleigh scattering and aerosols are also based on LOWTRAN7. Aerosols types within the model include boundary layer (rural, urban, and oceanic), tropospheric, and stratospheric (background stratospheric, aged volcanic, fresh volcanic,, and meteor dust). Aerosol optical depth, single scattering albedo and asymmetry factor can also be directly specified. These properties change with λ and the relative humidity of the cell they occupy.

Cloud optical properties are specified by the Henyey-Greenstein approximation or directly determined from Mie theory (Wiscombe, 1980). Cloud droplet effective radii are based on a drop size distribution specified by a gamma distribution. The phase function is discretized into 225 points with the majority of the points concentrated in the forward and backward peaks of the phase function. The effective radii range from 4 to 128 μm and any combination of these drops may distributed within the model domain. Additionally, ice particles (either Mie computed spherical ice or optical properties read in from a file) or drizzle can be mixed with droplets in a cell.

As noted above at step 3, it is necessary in the algorithm to determine the constituent that a photon has hit. A more computationally cost effective approach from the perspective of processing speed would be to skip this step and to combine the constituent optical properties to produce a composite phase function and single scattering albedo. However, the flexibility in being able to assign different cloud droplet size distributions, aerosol types and water vapor amounts for each cell independently could require a different phase function composite for every cell. The memory requirements for such storage can be computationally prohibitive.

In its most simplistic form, radiance estimated using the Monte Carlo method is computed by accumulating photon statistics within radiance bins that are centered on the viewing angle of a sensor. This method must assume a constant angular radiance across the aperture of the bin. Such an approach is computationally very expensive for all

but the largest angular bins as only a limited number of photons exiting the top of the atmosphere will fall within a given bin. In the SB3D approach, every scattering event within the model domain contributes a "radiance photon" to each sensor viewing angle. The weight, r_{wt}, of this photon is given by

$$r_{wt} = p_{wt} P(\mu,\phi;\mu',\phi') e^{-\tau} \qquad (1)$$

where p_{wt} is the incident photon weight, P is the particle phase function, (μ', ϕ') is direction of the incoming photon, (μ, ϕ) is the direction towards the sensor, and τ is the total optical depth from the scattering event to the sensor. This approach dramatically improves the statistics of the radiance estimate with the ratio of radiance photons to photons entering the top of the atmosphere ranging from 1:1 to more than 1000:1 depending on cloud optical thickness. Even with this improvement there still exists a problem of "hot spots" appearing in the radiance field that are caused by the large forward scattering peak of the Mie generated cloud phase function. To suppress this effect, P can be truncated and renormalized and τ and ω_0 rescaled. However, to maintain the integrity of the strong forward peak for single scattering solutions, the truncation and rescaling occur only after the initial scattering event for each photon.

By using the above method the directional tolerance of the radiance computation becomes irrelevant since there is no angular binning. For the case of cloud droplets with a phase function computed from Mie theory the effect of binning can smooth out significant angular details in back-scattered radiation. For instance, the difference between nadir reflectance and the reflectance a few degrees off nadir for a zero-degree solar zenith angle can be tens of percent.

The surface component of the model can be selected as either a simple Lambertian reflector or a ocean surface with 3-D waves. The Lambertian surface has an albedo that is spectrally varying depending on the type of surface selected (e.g. snow, sand, vegetation, lake, ocean). For the 3-D waves, photons are scattered or transmitted at the ocean surface based on the ocean-atmosphere refractive index interface. The wave facet angle and sea foam amount is calculated as a function of wind speed. For photons that penetrate the ocean surface, scattering and absorption are computed as a function of the optical properties of hydrosols and phytoplankton. Photons can be scattered back into the atmosphere providing a means for investigating the effects of the atmosphere on the retrieval of biological activity from satellite observations.

Computations are conducted in either a 3-D (3DM) or independent pixel approximation mode (IPM). For IPM, plane-parallel computations are essentially performed at each cell, and a photon is constrained to a horizontally homogeneous atmospheric column. After a photon enters the top of the model it remains in the same column as it traverses vertically. If the photon reaches the boundary of the column, it returns at the opposing boundary with the same trajectory, thus creating a cyclic boundary. By comparison, the photon can traverse horizontally for the 3DM computations thereby allowing it to encounter variations in optical depth and constituent microphysics in both the horizontal and vertical directions. Only at the edges of the model do the cyclic boundaries come into effect.

3. Accuracy

The accuracy of the flux estimates is proportional to the square root of the number of photons used in the simulation. For simple applications, the number of photons required is predetermined from Bernoulli probability based on estimates of the resulting flux and the desired level of random error. Since the goal of most simulations is to predict these fluxes, an exact estimate on the number of photons required prior to the initialization of the model run is not possible. After a model run is complete it is possible to estimate the error associated with the computed fluxes. However, as with most models the use of weighted photons in SB3D makes the estimate of the error difficult. When the values of these weights are a function of complex combinations (i.e. k-distribution weights, single scattering albedo and radiance weights), the estimate of the error using a Bernoulli type of approach is not feasible.

Another method towards estimating the error would be to repeat the model runs many times with a different initial random seed and to examine the combined statistics of the model outputs. However, the number of runs and number of photons processed per run must be defined and the eventual error estimate will be subject to these specified values. The simplest and probably most pragmatic approach is to judge the accuracy of the model output by monitoring its convergence and stability as the number of photons used in the simulation

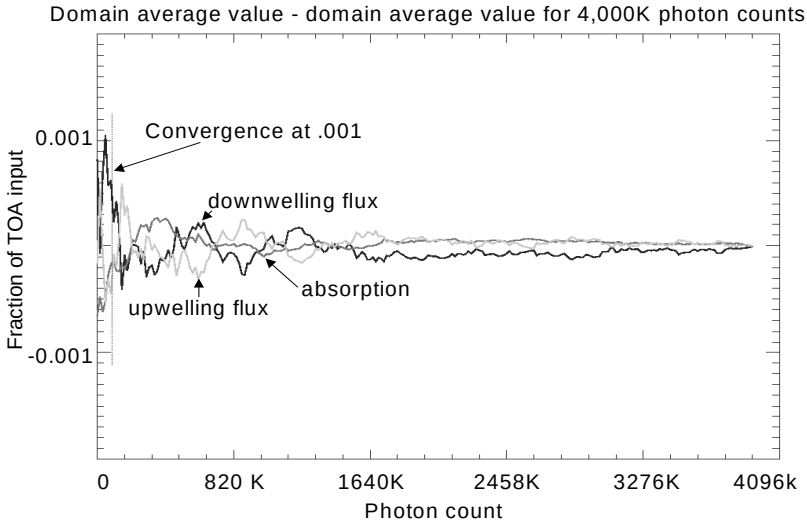

Fig. 2. Convergence of model domain averages. Dashed line points to first instance when all three domain averages have changed by less than .001 over three consecutive intervals of 16K photon counts.

increases. If checked at too short of an interval, however, convergence may be sensed at a local minimum and a model may be prematurely terminated. The determination of this interval is subjective and requires cognizance as to the quantity being estimated. For instance estimating domain average fluxes requires many fewer photons, and hence a shorter sampling interval than would be required for mapping the radiance of that same field.

In SB3D, a convergence criterion is applied, and the Monte Carlo process is terminated once a desired accuracy has been achieved. For domain average fluxes, the convergence is considered to occur when atmospheric absorption, transmission and reflectance as measured as a percent of the TOA input, all change by less than a given percentage *i* over three consecutive intervals of *j* photon counts. By testing the convergence on all three-domain averages rather than just one, the convergence is more stable. The stability of the convergence is demonstrated for the I3RC Landsat cloud field in Fig. 2. For mapping fluxes or radiance, the Monte Carlo process is terminated when convergence occurs for each atmospheric column measured individually. The convergence must occur simultaneously for all columns. For radiance, the same criterion is used but the relative change must not be more than the specified percentage *i* over three consecutive intervals of *j* photon counts.

4. Performance

A perennial decision in developing code is whether to optimize it for speed or for storage. For SB3D the philosophy was to develop a code as flexible as possible that could be used to address a variety of different research problems in 3-D radiative transfer rather than an operational tool. Thus, methods for storing many types of gas, aerosol and cloud optical properties necessitated algorithms that were not optimized for speed. Additionally, the desired output of the model can have a strong impact on processing speed. For instance, the model can be run in an irradiance mode only that skips the radiance computation in Eq. 1. For optically thick clouds, the effect of computing the radiance for each scattering event can slow the code by more than a factor of ten. Finally, the most important variable in comparing speed performance is the accuracy of the output. As noted, the accuracy is proportional to the square root of the number of photons used.

In the case of the I3RC Landsat cloud field, an estimated accuracy of .001 for the output means required less than 100K photons. However, if the processing is terminated at that point, values for individual pixels become virtually meaningless. An equivalent accuracy at the pixel level will require many orders of magnitude more photons, and an equivalent increase in processing time. Hence, a meaningful comparison of model performance requires a tighter definition as to what is being measured and a more universal methodology for arriving at how accuracy is determined. Processing times cited in the I3RC test cases for SB3D reflect the time required for achieving convergence at the .001 value for domain average fluxes. However,

the model runs were continued in order to improve the pixel level results.

5. Future Comparisons

From phase 1 of I3RC, it seemed that part of the intercomparison was to assess the performance of the various 3-D codes in terms of processing speed and memory usage by computing fluxes for the large Landsat scene. This measure could be useful if the codes were to be employed in an operational sense and if efficiency were a concern. It is also appropriate to use such a field for research on 3-D effects. However, since an intercomparison is not a "beauty contest" but more a means for testing the validity of codes it may be more informative to use fields (not necessarily realistic) that test robustness. This can be accomplished using small fields of simple geometric shapes with some extreme aspect ratios. Codes that are credible for these more "extreme" fields should be reliable for investigating 3-D effects on cloud fields that more closely resemble nature.

The natural evolution for phase 2 of I3RC would be to include experiments that incorporate additional atmospheric components such as Rayleigh scattering, aerosols and gases and to run the models in both spectral and broadband modes. At completion of phase 1 of I3RC, the validity of the 3-D transport portion of the various models should be established. It is difficult to imagine physical circumstances where the detrimental effects of introducing more complicated and realistic atmospheres into the codes could not be borne out by simply running the 3-D models in a plane-parallel mode and checking the results with more established plane-parallel atmospheric radiative transfer models. Still, unimaginable coding errors can make themselves known through additional comparisons, so performing phase 2 should still be a fruitful exercise.

Phase 2 could contain just 4 small fields (i.e. 16x16) of high spatial resolution with high aspect ratio clouds. The first could be a cloud field with absorbing droplets. A second could be the same field with the addition of a moderate absorbing gas that has a vertical distribution similar to water vapor. The third could again be the same cloud field but with a weakly absorbing gas and the addition of Rayleigh scattering. The final field could be the same as the previous but with the addition of an overlying ozone layer. These relatively simple fields should discern any deficiencies that may possibly be related to 3-D mechanisms in the models. For cloud fields, atmospheres and spectral bands relevant to a particular research area, more extensive independent comparisons could be performed in 1-D using established plane-parallel models.

References

Wiscombe, W. J., 1980: Improved Mie Scattering algorithms. *Appl. Opt*., **19**, 1505-1509.

Yang, S., P. J. Ricchiazzi, and C. Gautier ,in press: Modified correlated-k distribution methods for remote sensing applications. *J. Quant. Spectrosc. Radiat. Transfer.*